
ECHO


ECHO

ECHO…

- Is a set of Interoperable Registries:
  - Data Registry
  - Web Service Registry
  - Provider Registry

- Serves provider and resource information to clients so they can directly interact with peers. This is a hybrid model in peer-to-peer terminology

- Provides a set of open APIs through which communications take place to support machine-to-machine or human-machine interactions

- Provides flexible and accurate search, at the collection- and inventory-level, of resources

ECHO RFC

- This RFC provides information about the NASA EOS Clearinghouse (ECHO) system as an operational standard for metadata publishing, discovery, and access
- The RFC describes how data and service resources can be shared, promoted, and utilized by the wider Earth science enterprise using ECHO as a standard enterprise service
- RFC sections contain information about ECHO (releases 5.5-6.0) as a standard operational service designed and operated for Earth science enterprise (ESE) stakeholders including the science data systems built to support them.

RFC Overview

- Each section of the RFC introduces a function and how it works. Details about the function, including its operational maturity and significant level of support, are indicated through references to supporting online information.
- Taken together ECHO's publish, discover, and access functional capabilities serve the science community as a qualified "pattern" for NASA to apply to science infrastructures developed for sharing/using distributed resources

RFC

Presenting a case for ECHO as a qualified operational metadata system standard.

- **Background**
  - ECHO design concepts, key features, ECHO service concepts, advantages for ECHO partners

- **Publishing** – support metadata and information sharing
  - ECHO data partners
    - Ingest policies, mechanisms, API, management tool (PUMP)
  - ECHO service partners
    - Extended services, registry, interface, functionality, testbeds
- **Discovery** – support query and retrieval of shared metadata and information (ECHO content)
  - ECHO client partners
    - Client services, metadata model, query, collection & inventory search
    - Collection & granule results DTDs, query results visibility and presentation
- **Access** – provide means for ordering data resources discovered in ECHO
  - Orderable catalog items, order process,
    - Parts that make up an order, approach to creating and submitting an order

Characteristics of the ECHO Standard

- **Ease of Provider Participation:** Designed to be low cost and minimally intrusive, ECHO offers a set of standard ways for providers to interface with the system and a metadata exchange approach that accommodates existing providers and technology.
- **Data Model Consistency:** In order to ensure data model consistency and ease of use in the ESDIS community, ECHO currently uses the Bulk Metadata Generation Tool (BMGT) format developed by ECS and is working with ESDIS and ECS to deliver data to ECHO in a common format. To mitigate the risk of not being able to match all possible provider data models, ECHO provides adaptors to translate provider formats into ECHO formats upon ingest into ECHO.
- **Open System/Published APIs:** ECHO uses an open system approach and ensures that user interfaces fully address user/scientist needs by specifying and publishing domain APIs that accommodate independent ECHO clients. These APIs are independent of the underlying transport protocols used. Currently, ECHO is capable of communicating using Simple Object Access Protocol (SOAP) and Java Remote Method Invocation (RMI). Plans are in place to add a web services view of ECHO services. Other transport protocols can be added as necessary.
- **Extensibility of User Interface:** ECHO extensibility is assured by its component architecture that allows new capabilities and functions to be plugged in, modeling relationships between services/APIs/UIs and continued prototyping. ECHO has focused on middleware that enables many different types of user interfaces to access ECHO services via its APIs.
- **Evolutionary Development:** The ECHO system is developed in increments to allow for insight and feedback during the development cycle. Industry trends are followed and the use of commercial, off-the-shelf products is optimized.

Advantages of ECHO

- ECHO serves Earth science enterprise stakeholders by enabling data providers to register their scientific resources (data, services) in a shared, managed, accessible facility
- Enables Earth scientists to access information about these ‘published’ resources through application clients that interface with ECHO via industry standard APIs
- Using ECHO services scientists are better able to search, find, order, and make use of heterogeneous and distributed enterprise data assets

Key Advantages ECHO Offers the Earth Science Enterprise

- Uniform Access - portal to Earth science data and services
- Programming Interface – allows organizations to connect their own interfaces and programs
- Standard Messaging – presents a messaging interface based on XML
- Centralized Metadata Access – allows providers to share metadata and offload some search responsibilities

What Do Users Get From ECHO?

- Capability to search for data without knowing what or where the data are
- Distinct directory and inventory searches
- Ability to search extensive, common metadata model across all data sources, plus extended metadata for specific products
- Ability to search through large datasets, accurately picking out a few products among millions
- Catalog of services that is human or machine readable (client can read the description of web service and understand how to construct a request for the web service)
- Facility to combine multiple web services and data on behalf of the client/user.
- Direct access of online data, or order it on media
- Privileged access to restricted data
- 99% system availability
- Wide variety of tailored clients will be available shortly (some are search clients on subset of data, some are machine-2-machine clients that directly access data to feed modeling systems, …)

What Do Data/Service Providers Get?

- Automated tools to map metadata into ECHO systems catalog metadata (don't have to change data provider database)
- Metadata update tools for metadata changes
- Data and Service interoperability – service provider can offer services on someone else's data
- Ability to extend metadata model with product specific search attributes
- Flexibility to add algorithms best suited to their data (e.g. orbital search algorithm)
- Open machine to machine interface for search (APIs)
- Two-level access control (visibility and access) on data (services - future)
- Project supplied operations support for metadata mapping, metadata ingest, machine uptime, ...
- Reusable software to help jump start development of a tailored client for a specific user group with specific capabilities (e.g. search client for polar data, data access & analysis client, ...)
- Data/service provider can insert copies of their data/service metadata into ECHO, let the ECHO ops team worry about daily operations so that provider can concentrate on primary job of archive, distribution and services. Can develop and offer tailored client to their users.

Why use ECHO?

- Off load operational responsibility for search for data and services to the ECHO team by ingesting metadata into ECHO. System has 99% availability. Sysadmins on call 24/7. Fully redundant backup system should be operational within year.
- Data/service providers shielded from standards volatility
  - E.g. metadata automatically mapped through ECHO instead of converting all their stored metadata to new metadata models.
  - E.g. Interface (API) changes
- Data/service providers participate in new HQ initiatives painlessly
  - E.g. ECHO becomes new node in NSDI by addition of NSDI adapter software to ECHO system
  - E.g. ECHO becomes new node in GEOSS by addition of GEOSS adapter software to ECHO system
- Data and service providers can use ECHO's reusable software to develop a tailored client to access a subset of the ECHO holdings.
- ECHO is invisible to the users. Users only see the client and then the data or service access. Users don't know the metadata was held in ECHO or that the search was executed in ECHO.
- Data and service providers can concentrate their budget dollars on providing more and better data and services and fulfilling orders for data and services.

Current and Planned Use

- 18 Client Providers - providing views tailored to their community, application or modeling system
- 10 data providers - publishing resources for broader sharing
- REASoN CAN Winners:
  - NASA EOS Higher-Education Alliance (GMU) uses ECHO to find and access data in EOSDIS data pools and near-line storages for generating on-demand, value-added, educational and scientific products for users. Future plans to use ECHO for publishing services and for finding needed services for students to build geospatial processing models.
  - Invasive Species Forecasting System (Department of Interior, Department of Agriculture, various universities, State of Colorado and NASA) uses ECHO to access EOS data to drive models of invasion of non-native species so that government programs that control the invasion can be targeted to improve cost efficiency (programs to deal with Invasive Species cost US and state governments $137 Billion per year).

Current and Planned Use

- International Activities
  - The CEOP (Coordinated Enhanced Observing Period) coordinates research for the World Climate Research Programme's Global Energy and Water Cycle Experiment (GEWEX) Hydrometeorology Panel will study how to use ECHO as the inventory search engine to find and make CEOP data available to CEOP scientists via a custom CEOP client. There are also efforts to make NASA EOSDIS data available to CEOP scientists through OPeNDAP enabled clients and servers.
  - Planned interoperability with JAXA (Japan's Space Agency).
  - Studying interoperability with ESA and other international partners
- Various Proposed Efforts
- ESIP Federation (using ECHO Services Testbed)

LoE for Non-ECS, ECHO Data Partners

ORNL: 440 hours
- getting metadata into ECHO < 40 hours
- implementing a SOAP interface to accept orders: 400 hours

ASF: 600 hours

SEDAC: 168 hours
- Programmer: 30hrs/wk for 2 weeks = 60hrs
- Domain expert: 30hrs/wk for 2 weeks; 4hrs/wk for 2 weeks = 68hrs
- Metawriter: 20hrs/wk for 2 weeks = 40hrs

Stennis: Information not available.

Pull Partners:
Harvesting was on the ECHO workplan but is not considered core since we have the capability above. Given that a harvester needs to be tailored on the provider end for each provider, not knowing the variability and the complexity on the provider side, and the need to support update and delete capabilities to ensure quality of content of ECHO, I would say each harvester would take one person-year. I would recommend a real analysis.

LoE for ECHO Client Partners

- Variability in effort depends on a number of factors:
  - Degree of capability/complexity in UI functionality
  - Degree of re-use employed in building the UI
  - Familiarity of programmer with XML, etc
  - Timing of availability and maturity of ECHO feature to be exercised by client
- The real question is how much additional work must a client developer engage in to use ECHO over something like V0 or OGC CS, or perhaps their own local database
  - Major tasks a client developer must do to search a catalog:
    - Build the query – an ECHO query can be built with industry standard tools whereas V0 and OGC CS have no such support – I am not aware of toolkits for V0 or OGC similar to what we are proposing, but I don't know that there aren't any
    - Interpret the results – Again, an ECHO result can be parsed with industry standard tools – V0 and OGC CS do not have broad sets of industry supported tools
    - Submit the request – V0 requires sending and receiving text across network sockets, OGC CS is layered on top of HTTP, ECHO uses true SOAP – There is considerable broad industry support for SOAP, HTTP has broad support, V0 requires pretty much custom coding to transmit the request
    - Accessing Data – V0 and ECHO provide mechanisms to order data of about the same degree of complexity, both ECHO and OGC support direct URL links to data which are simple

LoE for ECHO Client Partners

- *Effort required to build a client has ranged from 4 person-weeks to 2 person-years.*
- *Building the interface to the ECHO APIs ranges from 2 person weeks to 16 person months depending on the degree to which APIs are used.*

Mercury EOS

Goal: Enable access to ORNL and other Enterprise holdings
Level of Effort: 2 FTEs for 1 year (includes interface to ECHO)

Data Validation User Interface (DVUI)

Goal: Allows MODIS scientists to rapidly find basis-of-validation granules and resultant granules from other sources
Level of Effort: 1 FTE for 1 year
Level of Effort to connect to ECHO: 1 FTE for 1 month

Power User Interface (PUI)

Goal: Allow scientists to order archived granules via only local granule ID
Level of Effort: 4 person weeks
Level of Effort to connect to ECHO: 1 FTE for 2 weeks

NASA Earth Observations (NEO)

Goal: Merging of Browse and ordering of data
Level of Effort: 2 man years (application),
Level of Effort to connect to ECHO: 4 man months

Invasive Species Forecasting Service

Goal: Predict invasion of non-native species in U.S.
Level of Effort: 1 FTE for 1 year
Level of Effort to connect to ECHO: 1 FTE for 1 month (ECHO)

AnnoTerra

Goal: Integration of content and metadata from heterogeneous ES systems using semantic web technology
Level of Effort: 1.5 person months (includes interface to ECHO)

MODIS Rapid Response System (RRS)

Goal: Allow RRS users to find and acquire datasets related to images on the RRS public interface that cannot be generated or stored
Level of Effort: 3 person months (includes interface to ECHO)

Warehouse Inventory Search Tool (WIST)

Goal: General Purpose access to data and services registered in ECHO
Level of Effort: 2 FTEs for 6 months
Level of Effort to connect to ECHO: 4 FTEs for 4 months

Earth-Sun Gateway (ESG)

Goal: Facilitate discovery and use of data and services by decision-makers, scientists and others needing access to NASA Earth Science (and Earth-Sun) data and services.
Level of Effort to connect to ECHO: 40 person-hours

NASA Earth Observations (NEO)

NASA MODIS Science Team (SSAI)

Motivation

- Target Audience: Informal educators (museums, science centers, etc.); non-NASA scientists; “science aware” general public.
- Goal: Initiated by the creators of the *Earth Observatory, NEO* will merge the capacity to quickly and easily browse EOS data with the ability to order/download the raw data. In the first system prototype, NEO’s web-based user interface will guide non-expert users in the discovery, exploration, and acquisition of EOS MODIS browse data and their underlying data sets regardless of the location of the archive at which the data are located.
- Current State: In development.
- Schedule for Completion: Fall 2005

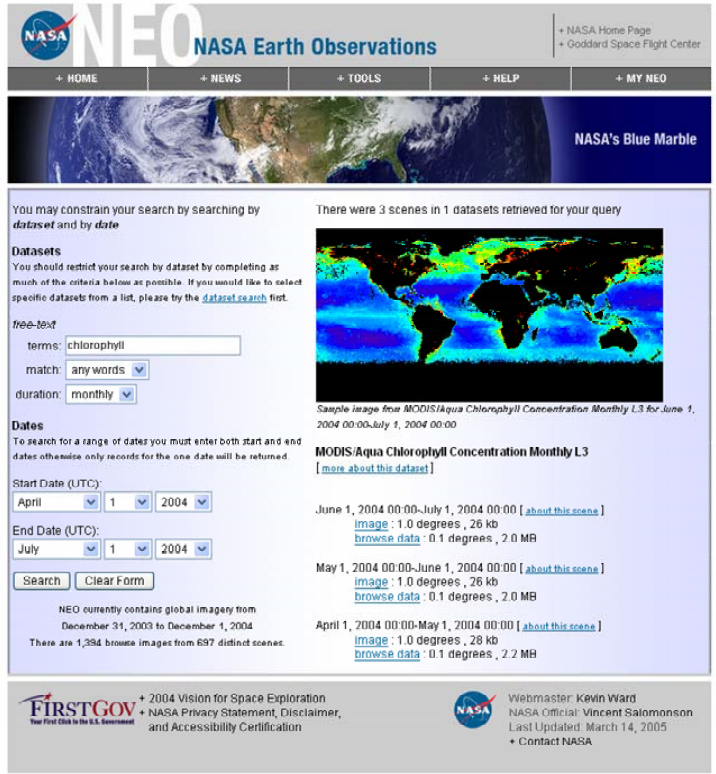

Results

- Once released, NEO will allow our audience to navigate numerous, complex EOS datasets and order data from multiple DAACs via a single interface.
- Level of Effort to Build Client: 2 person years
- How Much Code Re-Use:
  - utilizing the EchoTalk client library
- Level of effort to plug into ECHO: 4 person months

Impact of ECHO

- ECHO provides capabilities that NEO would be unable to duplicate. It would be impossible to provide a one-stop data shopping interface to our audience without ECHO as EOSDIS does not provide any machine-capable interfaces to query all of the data sources (i.e., DAACs). Without ECHO, users would not know which DAAC contained which dataset, and this virtually shuts-out our target audience from data access at this time.

AnnoTerra

SSAI

Motivation

- Target Audience: Semantic Web research community
- Goal: Develop a proof-of-concept system that automatically integrates content and metadata from heterogeneous NASA Earth science information systems (Earth Observatory, GCMD, ECHO) using Semantic Web technology.
- Current State: Prototype completed for 2003 Semantic Web Challenge; it is no longer publicly accessible.
- Schedule for Completion: Completed in Oct. 2003

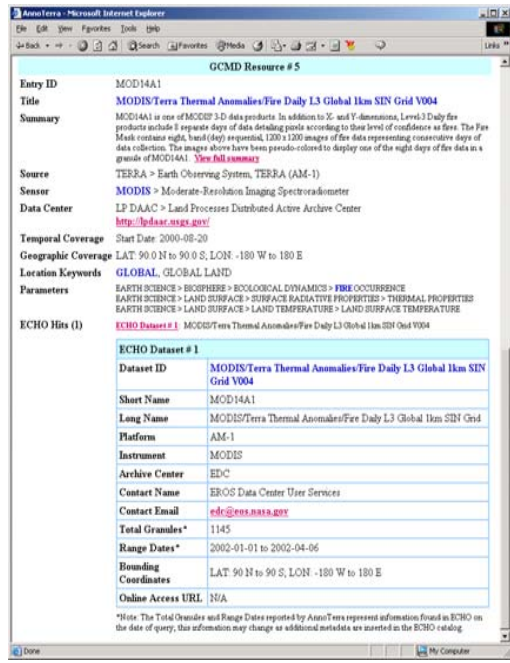

Results

Successful demonstration of the AnnoTerra prototype resulted in 3rd Prize in the 2003 Semantic Web Challenge (http://challenge.semanticweb.org/) and a publication in IEEE Intelligent Systems journal (Vol. 19, Number 3 May/June 2004).

- Level of Effort to Build Client: 1.5 person months
- How Much Code Re-Use: 90% - used Java library from ECHO Operations Team
- Level of effort to plug into ECHO: N/A – included in Level of Effort to Build Client

Impact of ECHO

- Before ECHO … Without ECHO, the AnnoTerra prototype would not have been able to incorporate the identification of inventory-level resources. Key words extracted from the Earth Observatory news feeds would have been mapped only to catalog-level resources in GCMD. The developers of AnnoTerra tried to find other sources of machine-accessible inventory-level metadata. In 2003, only a few others could be found, and none of them provided their native data model in an appropriate format for machine use.
- After ECHO … ECHO enabled the incorporation of inventory-level resources in the AnnoTerra prototype, which its developers believe was a key factor for success in the Semantic Web Challenge.

Invasive Species Forecasting Service (M2M)

NASA CAN, SESDA

Motivation

- Target Audience: Modeling System for modeling distribution and likely habitats for the spread of invasive species. This is a Machine-to-Machine interface where the “user” is the Invasive Species Forecasting System Data Component.
- Goal: Predict invasion of non-native species to drive decisions on where to target efforts to study and control the invasion. Currently the US gov’t spends an estimate $133Billion/year to deal with invasive species.
- Current State: In Development.
- Schedule for Completion: Dec ‘05 ECHO client Alpha release

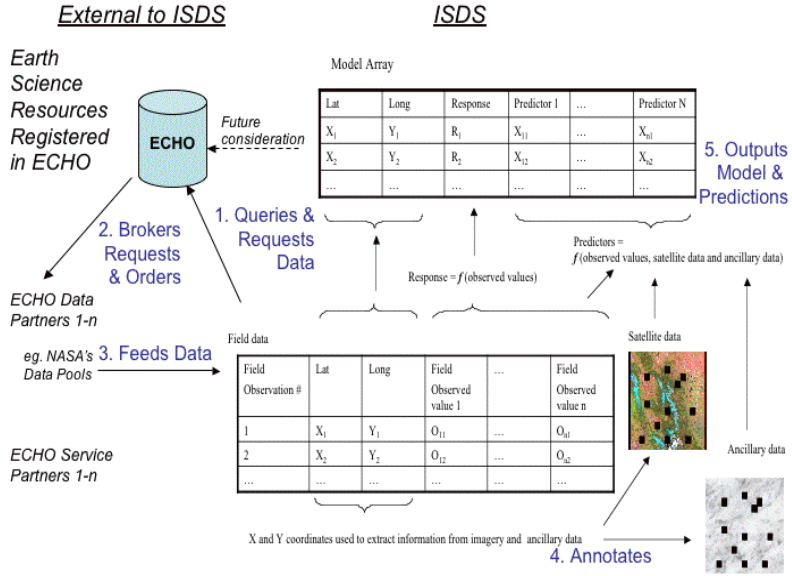

Results

(description of expected results)

- Level of Effort to Build Client: 12 man months
- How Much Code Re-Use: 50% estimate - w/ use of JavaTalk lib and collaborate with existing client development teams
- Level of effort to plug into ECHO: minimal, 1 month

Impact of ECHO

- Before ECHO….
  - Define custom Earth science metadata model
  - Learn & interface w/ various DAAC custom client interfaces
  - Develop some common access & query mechanism compatable across DAACs
  - Stage data manually for use in ISFS
- After ECHO….
  - Standard Earth science metadata model defined, populated and kept current in a standard and central repository
  - Provides one common interface to repository, with multple machine interfaces (Java, web services, etc.)
  - Online spatial & temporal query of metadata, with realtime, online data download to ISFS server.
  - Easy metadata record and web services syncronization w/ GCMD
  - ISFS developed data products may be registered in ECHO for use by collaborating scientists.

DVUI - Data Validation User Interface

GSFC Professional Intern Program

Motivation

- Developed to support MODIS Land Validation Team ~10 users. Can extend to any comparable data validation efforts ~100 users.
- Provides scientists a custom means for rapidly finding basis-of-validation data granules and then find coincident granules from other sources. Enables users to select granules that:
  - Are as cloud free as possible
  - Overlap as much as possible (to reduce the number of granules to just those that are relevant)
  - Are close in time so that changes that occur over time do not influence the validation results
- Provide a pre-selected view of metadata from 3 years data represented by 37 datasets over 31 discrete validation sites updated daily for streamlined access to validation granules.

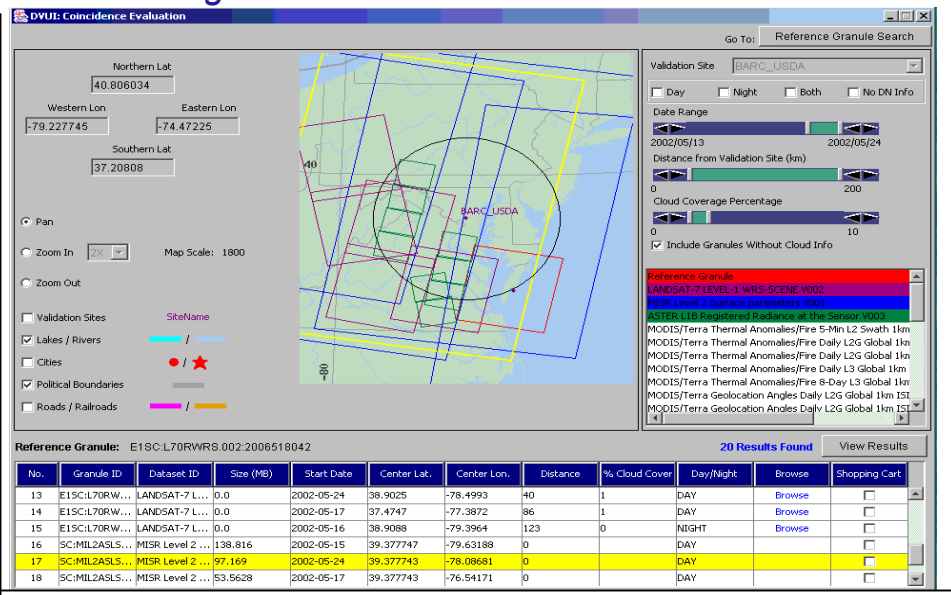

Results

- Searches performed in minutes compared to hours or days
- Level of Effort to Build Application:
  - GUI: 1 person year ("fresh-out" in GSFC's Intern program)
  - Harvester: 1 person year
- Code Re-Use: client 50% (GLISmapper - USGS, StarDOM – UMD, Java Examples - Sun); harvester 10% (path/row and lat/lon algorithm - EDG, day/night flag algorithm - ECS, SOAP perl proxy - ECHO)
- Bulk and update mode metadata harvester implemented to synchronizes DVUI pre-query content with ECHO
- Level of effort to plug into ECHO: 7 person-weeks
  - Interface to ECHO SOAP API straight forward: 2 person-weeks
  - DTD-to-ECHO DTD query straight forward: 1 person-week
  - Query optimization/QA for metadata harvesting: 1 person-month

Impact of ECHO

- Pre-ECHO systems could not generate metadata views for a single instance (not considering refresh) of the 3 x 37 x 31 query result set on a daily basis. DVUI was impossible with old architecture systems.
- ECHO enables DVUI metadata harvesting so data validation teams can access up-to-date data on a daily basis per their requirements.
- Before ECHO, cross-provider machine-to-machine interaction and metadata harvesting were not possible in EOSDIS.
- ECHO reduces data access time from days to hours by enabling functions previously supported in 4 different user interfaces to be accomplished in in a single user interface.
- ECHO enables alternate data access paradigms to find data from distributed archives.

Earth-Sun System Gateway (ESG)

NASA Geospatial Information Office, Synergy

Motivation

- Target Audience: Decision-makers, scientists and others needing access to NASA Earth Science (and Earth-Sun) data and services.
- Goal: Facilitate discovery and use of data and services
- Tailored to showcase NASA applications of national interest, but not limited to these applications.
- New capabilities easily adapted through standards conforming interfaces.
- Current State: beta test; transitioning to public release
- Schedule for first iteration:
  - August 2005: Public release
  - Fall 2005: populating database, publicizing, and gathering further input

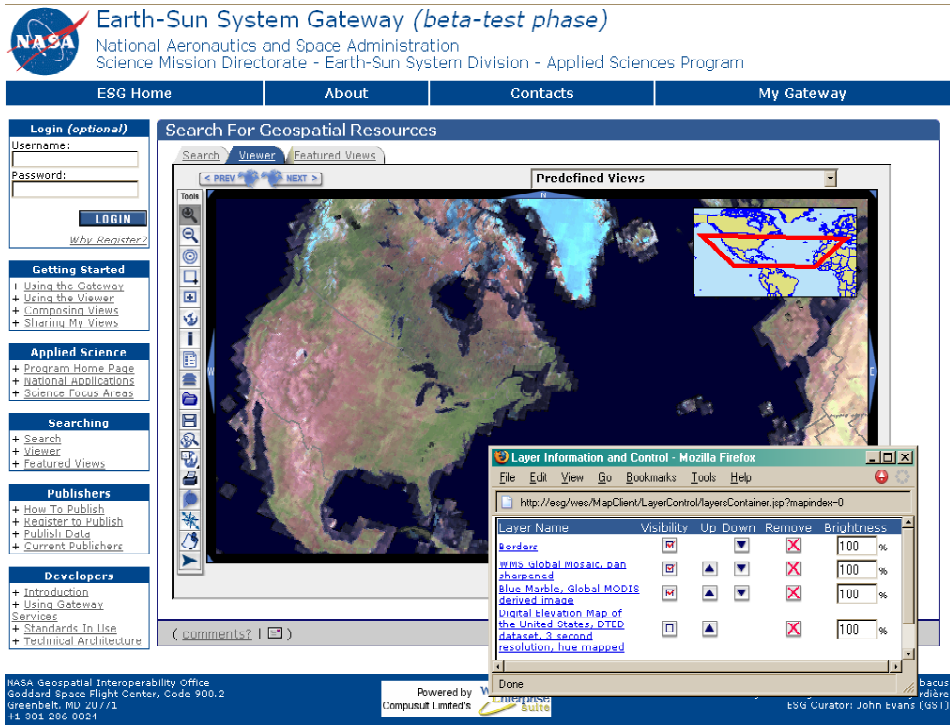

Results

- Access EOS data on ECS Data Pools through NSDI adaptor to ECHO catalog (Synergy V)
- Online portal: OGC / Z39.50 catalog client, WMS viewer, and ebRIM search/publish access, all in an HTML environment
- Portal users register resources (data, services, documents &c.); other users can find them, reuse them, and invoke Web services to visualize data and models.
- Portal also harvests NSDI metadata (~200 catalog servers)
- Catalog service (usable by other OGC Catalog clients)
- Gazetteer service (usable by other OGC Gazetteer clients)
- Ability to provide specialized “viewpoints” for each National Application and Science Focus Area
- WCS client support and 3D viewer links (Synergy V)
- Level of effort to plug into ECHO: 40 person-hours

Impact

- Showcased outputs of Synergy V interoperability work
- Data from CADRE, MODIS AIRNow, SERVIR, CQUEST, and other NASA application collaboratives now accessible through ESG; working on full WMS integration.
- In use as a compendium of online resources of types related to Earth Science and the Earth-Sun System
- Upon deployment of newly developed OGC Web Services capabilities for DAAC Datapools, DAACs can make holdings available to ESG through the ECHO catalog.

MODIS Rapid Response System

NASA NRA and MODIS, SSAI contract

Motivation

- Target Audience: Users of the Rapid Response System (RRS) public web interface – more than 400,000 unique visitors were identified for the RRS in 2004. RRS users are distributed evenly among 3 groups: scientists and applications users, media organizations, and the general public (http://rapidfire.sci.gsfc.nasa.gov/).
- Goal: Allow RRS users to quickly find and acquire datasets that are related to the images posted on the RRS public interface (more than 8,000 per day) that cannot be generated and/or stored in the RRS due to resource limitations.
- Current State: In Development – Requirements Phase
- Schedule for Completion: Prototype by Nov. 2005

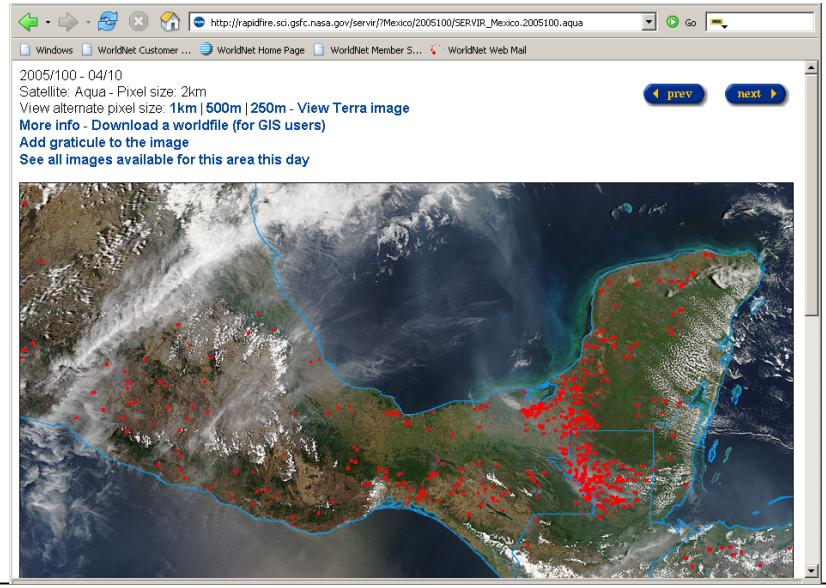

Results

When the "Direct Connect" interface is completed, RRS users will be able to search and order related datasets from multiple data centers without having to manually navigate to other data access tools (e.g. EDG, WHOM) and manually transfer their detailed search parameters.

- Level of Effort to Build Client: 3 person months (estimate)
- How Much Code Re-Use: 75% - using Java library and example client from ECHO Operations Team
- Level of effort to plug into ECHO: N/A – included in Level of Effort to Build Client

Impact of ECHO

- Before ECHO … RRS staff spend a lot of time answering questions from users about how they can get access to "data". This is because the RRS public interface is primarily distributing image products. While RRS is continually adding more products to support users needs, it will never be able to provide all of the data products that its users need. Many RRS users do not have the knowledge and/or resources to go somewhere else to find and acquire data.
- After ECHO … Within a few "clicks", users can find out about datasets available from a variety of distributed data centers that are related to RRS products. If they find items that meet their needs, they can order them directly from the RRS without having to get instructions from RRS staff and go to other websites.

WIS^{3}R^{3}D

NSIDC

Motivation

- Target Audience: Public
- Goal: Allow users to order a wide variety of data and have it all delivered in a uniform grid.
- Current state: ECS data are generally all in HDF format – but the coverage, spatial type, and resolution of disparate datasets can vary quite a bit. Processing the data into a uniform grid to facilitate comparison is a significant effort for each user. Centralizing that effort frees user resources and time to do Science.

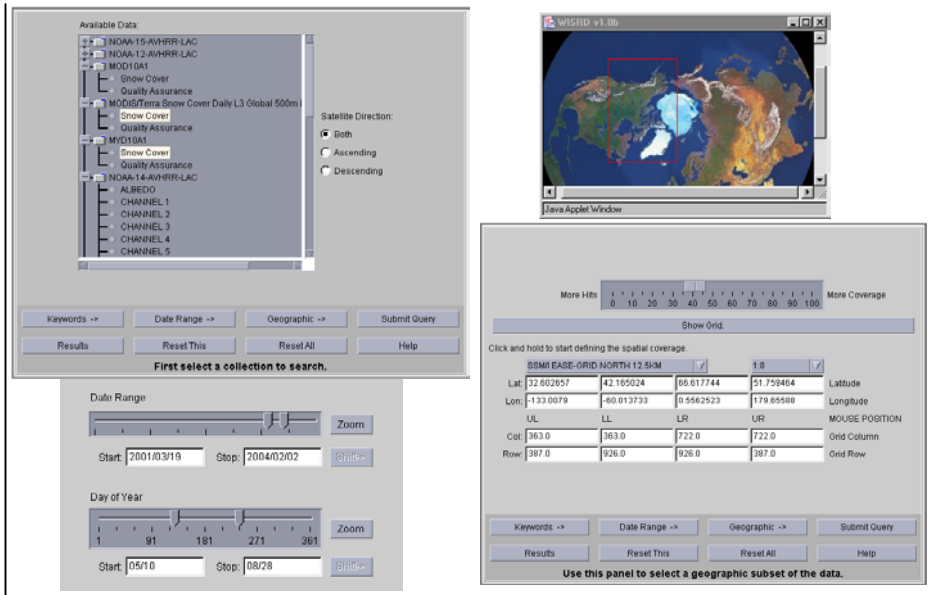

Results

- The WISRD interface allows users to choose their preferred grid and select a subset of that grid. That area is then used for the spatial search.
- WISRD backend processing subsets and/or stitches the data, grids or re-grids it to the area selected by the user, and outputs the data in the users preferred format.
- Level of Effort to Build Client:
  - The interface reused components from similar interfaces and took approximately 1 person month to develop.
  - Backend processing has to be developed per dataset and LoE can vary quite a bit.
- Level of effort to plug into ECHO: High, due to the need for introducing partner-provided orbital search model code.

Impact of ECHO

- Before ECHO:
  - There was no easy way to search or order ECS data holdings from an external client
  - Spatial search for orbital data near the poles returned many false positive results. In general the spatial search capabilities were costly, inflexible, and of mixed accuracy.
- With ECHO:
  - There are alternate ways for ordering data holdings from ECS and other providers
  - ECHO supports multiple spatial search mechanisms for increased accuracy in results sets
  - ECHO has the flexibility to continually improve while keeping configuration overhead low

WIST - Warehouse Inventory Search Tool

NASA/ESDIS

Motivation

- Target Audience: General Purpose access to data and services registered in ECHO
- Goal: General all purpose tool that fosters cross provider/discipline searches and order.
- Current State: Phase 1 development in Test.
- Schedule for Completion (of all phases):
  - WIST with ECHO 6.0: June 2005
  - WIST with ECHO 7.0: November 2005
  - WIST with ECHO 8.0: May 2006
  - Fully functional WIST operational: Oct 2006

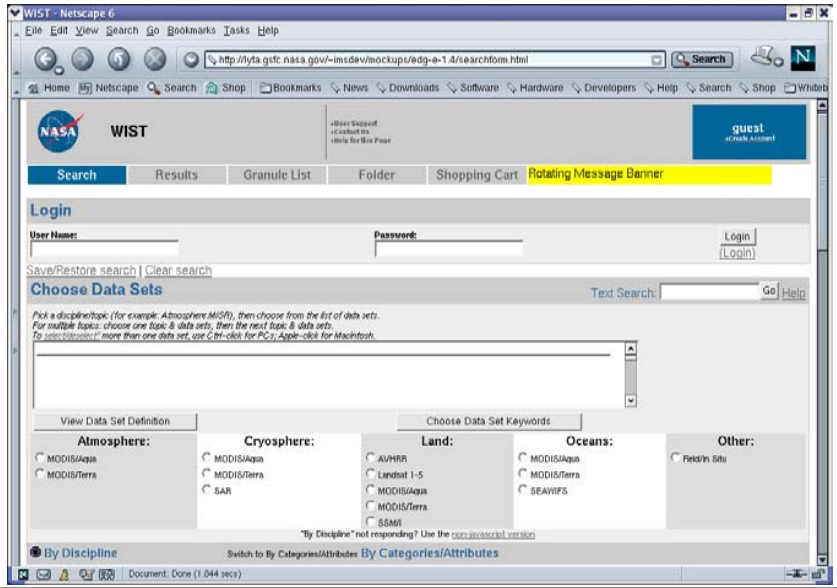

Results

- Client simplified significantly compared to its predecessor EDG - reduced maintenance costs
  - Client based distributed search eliminated.
  - Client no longer subject to providers being unavailable or slow to respond.
  - Client can more quickly implement new API features – not restricted by lagging data providers.
  - Need for provider specific extensions eliminated.
- Increased usability and some search performance.
- Level of Effort to Build Client: 2 FTE for 6 months.
- Code Reuse: ~70% reused from EDG
- LOE to plug into ECHO: 4 FTE for 3 months

Impact of ECHO

- Eliminates need for distributed searches thereby improving reliability and performance.
- Provider down time does not impact users
- API extensions can be rolled out as needed without having to coordinate being in lock step with numerous other groups
- Provider specific details insulated from clients.

Mercury EOS

Oak Ridge National Lab (ORNL), DOE

Motivation

Target Audience: Public, Science Users
Goal: Enable access to ORNL and other Enterprise holdings with the added value of additional free text search capabilities.
Current state: Developed for ORNL data.

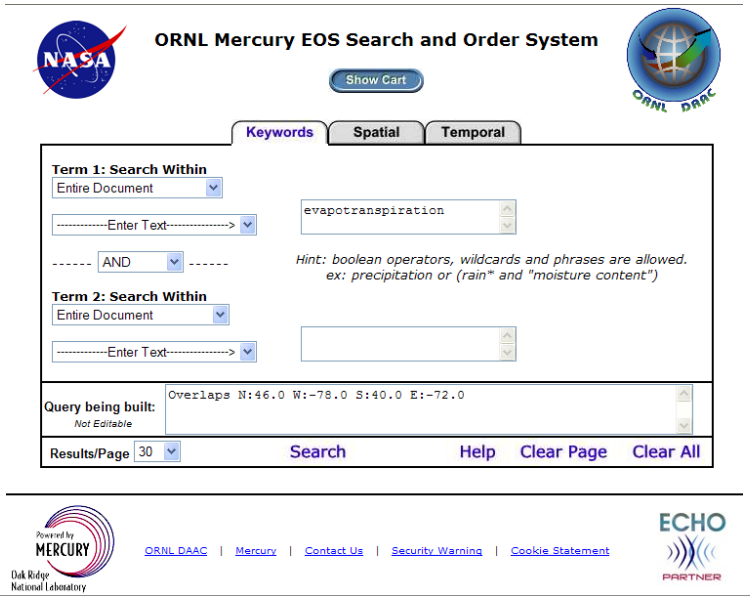

Results

Mercury-EOS harvests collection level metadata from ECHO nightly and augments with additional free text information. It's free text search engine helps users find new datasets that they are not currently familiar with. It places orders through ECHO.
Level of Effort to Build Client: 2 FTEs for 1 year.
How Much Code Re-Use: Shopping cart has been used on a MODIS prototype Web site
Level of effort to plug into ECHO: Included in above, client was designed for ECHO

Impact of ECHO

Provides open API for search, order, and order status. These were completely new capabilities. ECHO uses more modern Web technology than the V0 system.
Centralization of inventory with appropriate resources is a better performance and availability model than distributed queries.
As an ECHO provider, will allow discontinuation of our V0 server and old hardware maintained for it.

Power User Interface

NASA/ESDIS & ORNL

Motivation

- Target Audience: Science Teams (Knowledgeable Users, 10s of users)
- Goal: Allow science team members to order archived granules with just the local granule identifier in an automated fashion
- Current state: Operational

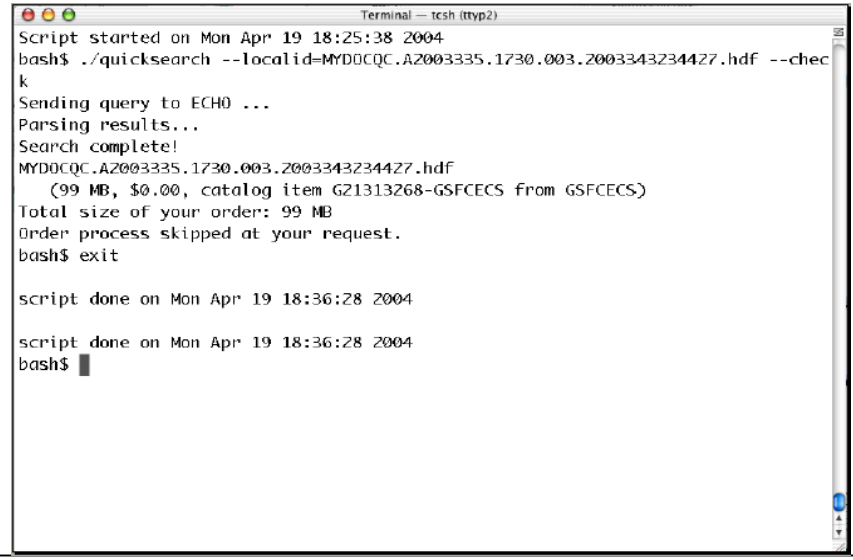

Results

- Quick Search is a command line tool that takes a list of granules and orders them through ECHO
  - All granules are ordered using the same options
  - User must only provide the list of granules, and can then walk away
  - User time is cut to almost nothing (minutes)
- Level of Effort to Build Client: 4 person weeks.
- Code Re-Use:
  - PUI application is all innovated code using perl support modules (CPAN installation compliant)
  - Web-based mode interfaced with ORNL “Shopping cart server” (interoperable use)
- Level of effort to plug into ECHO: 1 FTE for 2 weeks

Impact of ECHO

- Before ECHO, there was no efficient mechanism for bulk order from the ECS archives. Science team members had to order granules through EDG graphic user interface
  - Each granule requires no less than 7 clicks to order
  - No automated way to order a list of granules
  - Hours spent searching for and ordering granules
- This has streamlined the process to take seconds.